

Computer Systems for AI-programmers
baiCSAI3, Spring 2010
Lab Assignment: Code Optimization
Assigned: Week 9, Due: Week 10
Wednesday March 10, 24:00

Arnoud Visser (A.Visser@uva.nl) is the lead person for this assignment.

1 Introduction

This assignment deals with optimizing memory intensive code. Image processing and graphics offers many examples of functions that can benefit from optimization. In this lab, we will consider a single graphics operation: `rotate`, which rotates an image over 90° .

For this lab, we will consider an image to be represented as a two-dimensional matrix M , where $M_{i,j}$ denotes the value of (i, j) th pixel of M . Pixel values are triples of red, green, and blue (RGB) values. We will only consider square images. Let N denote the number of rows (or columns) of an image. Rows and columns are numbered, in C-style, from 0 to $N - 1$.

Given this representation, the `rotate` operation can be implemented quite simply as the combination of the following two matrix operations:

- *Transpose*: For each (i, j) pair, $M_{i,j}$ and $M_{j,i}$ are interchanged.
- *Exchange rows*: Row i is exchanged with row $N - 1 - i$.

2 Logistics

You may work in a group of up to two people in solving the problems for this assignment. The only “hand-in” will be electronic. Any clarifications and revisions to the assignment will be mailed to your student.uva.nl adress.

3 Hand Out Instructions

Start by copying `perflab-handout.tgz` to a protected directory in which you plan to do your work. Then give the command: `tar zxvf perflab-handout.tgz`. This will cause a number of files to be unpacked into a directory. The primary C-file you will be modifying is `rotate-versions.c`. In this file you can define multiple versions of the function `rotate()`. To test the performance of a version, you register this version to a benchmark-list in `kernel.c`, the second file that you have to modify. To be able to register your versions of the `rotate()` function, the name and description of each version has to be known, which you can define in `rotate-versions.h`. The main routine can be found in `driver.c`, which is a driver program that allows you to evaluate the performance of your solutions. Use the command `make driver` to generate the driver code and run it with the command `./driver`.

Looking at the file `kernel.c` you'll notice a C structure `team` into which you should insert the requested identifying information about the one or two individuals comprising your programming team. **Do this right away so you don't forget.**

4 Implementation Overview

Data Structures

The core data structure deals with image representation. A `pixel` is a struct as shown below:

```
typedef struct {
    unsigned short red;    /* R value */
    unsigned short green; /* G value */
    unsigned short blue;  /* B value */
} pixel;
```

As can be seen, RGB values have 16-bit representations (“16-bit color”). An image `I` is represented as a one-dimensional array of `pixels`, where the (i, j) th pixel is `I[RIDX(i, j, n)]`. Here `n` is the dimension of the image matrix, and `RIDX` is a macro defined as follows:

```
#define RIDX(i, j, n) ((i)*(n)+(j))
```

See the file `defs.h` for this code.

Your task is to rewrite the `rotate` code to make it run as fast as possible using techniques like code motion, loop unrolling and blocking.

See the file `rotate-versions.c` for this code.

Performance measures

Our main performance measure is *CPE* or *Cycles per Element*. If a function takes C cycles to run for an image of size $N \times N$, the CPE value is C/N^2 . Table 1 summarizes the performance of the naive implementations shown above and compares it against an optimized implementation. Performance is shown for 5 different values of N . All measurements were made on Pentium Dual E2180 machines.

The ratio (speedup) of the optimized implementation over a Baseline one will constitute a *score* of your implementation. To summarize the overall effect over different values of N , we will compute the *geometric mean* of the results for these 5 values. That is, if the measured speedups for $N = \{256, 512, 1024, 1536, 2048\}$ are R_{256} , R_{512} , R_{1024} , R_{1536} , and R_{2048} then we compute the overall performance as

$$R = \sqrt[5]{R_{256} \times R_{512} \times R_{1024} \times R_{1536} \times R_{2048}}$$

Assumptions

To make life easier, you can assume that N is a multiple of 32. Your code must run correctly for all such values of N , but we will measure its performance only for the 5 values shown in Table 1.

5 Infrastructure

We have provided support code to help you test the correctness of your implementations and measure their performance. This section describes how to use this infrastructure. The exact details of each part of the assignment is described in the following section.

Note: The only source files you will be modifying are `kernels.c`, `rotate-versions.c`, `rotate-versions.h`.

Versioning

You will be writing many versions of the `rotate` routine. To help you compare the performance of all the different versions you've written, we provide a way of "registering" functions.

For example, the file `kernels.c` that we have provided you contains the following function:

```
void register_rotate_functions() {
    add_rotate_function(&rotate, ROTATE_DESCR);
}
```

This function contains one or more calls to `add_rotate_function`. In the above example, `add_rotate_function` registers the function `rotate` along with a string `ROTATE_DESCR` which is an ASCII description of what the function does. See the file `rotate-versions.h` to see how to create the string descriptions. This string can be at most 256 characters long.

Driver

The source code you will write will be linked with object code that we supply into a `driver` binary. To create this binary, you will need to execute the command

```
make driver
```

You will need to re-make `driver` each time you change the code. To test your implementations, you can then run the command:

```
./driver
```

The `driver` can be run in four different modes:

- *Default mode*, in which all versions of your implementation are run.

Test case	1	2	3	4	5		
Method	N	256	512	1024	1536	2048	Geom. Mean
Optimized <code>rotate</code> (CPE)		12.7	16.5	25.1	18.7	66.9	
Baseline <code>rotate</code> (CPE)		17.2	31.3	103.6	98.1	107.4	
Speedup (naive/opt)		1.4	1.9	4.1	5.6	1.6	2.5

Table 1: CPE and Ratio for Optimized vs Baseline Implementations

- *Autograder mode*, in which only the main `rotate()` function are run. This is the mode we will run in when we use the driver to grade your handin.
- *File mode*, in which only versions that are mentioned in an input file are run.
- *Dump mode*, in which a one-line description of each version is dumped to a text file. You can then edit this text file to keep only those versions that you'd like to test using the *file mode*. You can specify whether to quit after dumping the file or if your implementations are to be run.

If run without any arguments, `driver` will run all of your versions (*default mode*). The naive version provided to you is very slow, and it will take several minutes to test. You will like to test only versions of your own very soon. When you made progress, copy and rename the function, to be sure that overwrite good working versions. Document your progress in a `labbook.txt`. Other modes and options can be specified by command-line arguments to `driver`, as listed below:

- g : Run only the final `rotate()` function (*autograder mode*).
- f <funcfile> : Execute only those versions specified in <funcfile> (*file mode*).
- d <dumpfile> : Dump the names of all versions to a dump file called <dumpfile>, *one line* to a version (*dump mode*).
- q : Quit after dumping version names to a dump file. To be used in tandem with -d. For example, to quit immediately after printing the dump file, type `./driver -qd dumpfile`.
- h : Print the command line usage.

Team Information

Important: Before you start, you should fill in the struct in `kernels.c` with information about your team (group name, team member names and email addresses). This information is just like the one for the Data Lab.

6 Assignment Details

Some advice. Focus on optimizing the inner loop (the code that gets repeatedly executed in a loop) using the optimization tricks covered in class. Look first at the algorithm, then at the details. Look at the assembly code for the `rotate` details.

Coding Rules

You may write any code you want, as long as it satisfies the following:

- You don't need `gcc-2.05.3` anymore. You should comment out the two packages needed in the previous assignment:

```
## Computer Systems:
#binutils-2.17
#gcc-2.95.3
```

- It must be written in C. You may not use any embedded assembly language statements.

- It must not interfere with the time measurement mechanism. You will also be penalized if your code prints any extraneous information.
- The tricks described in the website 'Achieving Greater Parallelism with SIMD Instructions'¹ are allowed, but are seen as bonus material (after the tricks described in the book).

You can only modify code in `kernel.c`, `rotate-versions.c`, `rotate-versions.h`. You are allowed to define macros, additional global variables, and other procedures in these files.

Evaluation

The score for each will be based on the following:

- Correctness: You will get NO CREDIT for buggy code that causes the driver to complain! This includes code that correctly operates on the test sizes, but incorrectly on image matrices of other sizes. As mentioned earlier, you may assume that the image dimension is a multiple of 32.
- CPE: You will get full credit for your implementations of `rotate` if they are correct and achieve mean speedups above 5x the baseline. You will get partial credit for a correct implementation that does better than 2x the baseline.
- Innovation: You will get extra credit for your achievements as documented in the `labbook.txt`.

7 Hand In Instructions

When you have completed the lab, you will hand in one file, `teamname-version-perflab.tgz`, that contains your solution. Here is how to hand in your solution:

- Make sure you have included your identifying information in the team struct in `kernel.c`.
- Make sure that the `rotate()` function correspond to your fastest implementation, as this is the only function that will be tested when we use the driver to grade your assignment.
- Remove any extraneous print statements.
- Describe your achievements in the file `labbook.txt`
- To handin your tar file, type:

```
make handin TEAM=teamname
```

where `teamname` is the team name defined in your team struct.

- After the handin, if you discover a mistake and want to submit a revised copy, type

```
make handin TEAM=teamname VERSION=2
```

Keep incrementing the version number with each submission.

Good luck!

¹<http://csapp.cs.cmu.edu/public/waside/waside-simd.pdf>